

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Patent Application

Applicant(s): Y. Diao et al.
Docket No.: YOR920030088US1
Serial No.: 10/648,179
Filing Date: August 26, 2003
Group: 2123
Examiner: Juan Carlos Ochoa

Title: Methods and Systems for Model-Based
Management Using Abstract Models

APPEAL BRIEF

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Applicants (hereinafter referred to as "Appellants") hereby appeal the final rejection of claims 1-33 of the above-identified application.

REAL PARTY IN INTEREST

The present application is assigned to International Business Machines Corporation, as evidenced by an assignment recorded August 26, 2003 in the U.S. Patent and Trademark Office at Reel 14442, Frame 276. The assignee, International Business Machines Corporation, is the real party in interest.

RELATED APPEALS AND INTERFERENCES

There are no known related appeals or interferences.

STATUS OF CLAIMS

Claims 1-33 stand finally rejected under 35 U.S.C. §103(a).

STATUS OF AMENDMENTS

There has been no amendment filed subsequent to the final rejection.

SUMMARY OF CLAIMED SUBJECT MATTER

Independent claim 1 recites a method of constructing a model representative of a resource for use in managing a service associated with the resource, comprising the steps of associating a resource abstract model with the resource, wherein the resource abstract model is configured to automatically determine a set of resource metrics to be used to construct a model representative of the resource such that a reduced set of resource metrics is considered; and constructing the model representative of the resource based on the reduced set of resource metrics obtained in accordance with the resource abstract model. The present specification provides an illustrative embodiment of the elements of claim 1 at page 12, line 15 through page 13, line 2.

More particularly, FIG. 4 of the present application depicts an overall process for automated construction and exploitation of quantitative models using abstract models, according to an embodiment of the present invention. Step 405 illustrates the step of associating a resource abstract model with the resource. In step 405, one or more RAMs are constructed by resource experts for that type of resource. That is, RAMs are constructed by one or more subject matter experts for the measurement source. For example, a RAM for a web server access log might be constructed by a web server design and/or support team. The model builder uses the RAM to select metrics based on logical relationships.

The resource abstract model is configured to automatically determine a set of resource metrics to be used to construct a model representative of the resource such that a reduced set of resource metrics is considered. For example, there are many measures of paging (e.g., page-ins, page-outs, faults, page cache hits). The RAM specifies that these are related metrics so that the model builder knows that not all of them may be needed in building the model. Using the RAM

reduces the model dimension, which has at least three important consequences: (1) less computation is required to construct the model; (2) less data is collected since fewer parameters are estimated and so the model can adapt more quickly; and (3) since there are fewer parameters, the model tends to be more robust to noise.

The step of constructing the model representative of the resource based on the reduced set of resource metrics obtained in accordance with the resource abstract model is illustrated in FIG. 5, a flow diagram illustrating a process for building a system model using abstract models, according to an embodiment of the present invention. This process may be considered an illustrative explanation of step 410 of FIG. 4. In step 505, metrics are obtained from the service level agreement (SLA). The obtained metrics are stored in history file 412 (FIG. 4). In step 510, the configuration database is queried to get the service topology for the service of interest. In step 515, the RAMs are obtained from the resources identified in the service topology. In step 520, a subset of the metrics available in the identified resources is identified based on considerations such as whether they relate to the features used in the service for the service level of interest. In step 525, metric data is collected from the resources using management-agent interfaces. The collected metric data is stored in history file 412. In step 530, techniques such as principle components analysis and stepwise regression are used to select the metrics to use in the system model. Data is read from history file 412 in step 530. In step 535, the model itself is constructed (e.g., by estimating the constants in a linear regression model). (Specification, page 12, line 15 through page 13, line 27).

Independent claim 11 is an apparatus claim having similar elements as the above-described claim 1. The apparatus comprises a processor and memory arrangement. FIG. 8 shows an illustrative hardware implementation a computing system in accordance with which one or more components/methodologies of a management system may be implemented. As shown, the computer system may be implemented in accordance with a processor 802, a memory 804, I/O devices 806, and a network interface 808, coupled via a computer bus 810 or alternate connection arrangement. (Specification, page 15, line 23 through page 26, line 15). Therefore, FIG. 8 in conjunction with the features described above with regard to FIGS. 4 and 5 serve as an illustration of the language of claim 11.

Independent claim 20 is an article of manufacture claim having similar elements as the above-described claim 1. The article of manufacture comprises a machine readable medium. The present specification provides an illustrative embodiment of the elements of claim 20 at page 17, lines 6-10. More particularly, as explained therein, software components including instructions or code for performing the methodologies described herein may be stored in one or more of the associated memory devices (e.g., read only memory, fixed or removable memory) and, when ready to be utilized, loaded in part or in whole (e.g., into random access memory) and executed by a CPU. (Specification, page 17, lines 6-10). Therefore, the present specification at page 17, lines 6-10, in conjunction with the features described above with regard to FIGS. 4 and 5 serve as an illustration of the language of claim 11.

Independent claim 27, having similar elements as the above-described claim 1, recites the additional steps of deploying one or more resource abstract models in association with one or more resources, and using the one or more constructed models to manage the one or more resources. The present specification provides an illustrative embodiment of the elements of claim 27 at page 14, lines 16-21. More particularly, as depicted in FIG. 7, a flow diagram illustrates a process for how service providers prepare and distribute information to enable automated operation, according to an embodiment of the present invention. In step 705, the service level agreement is constructed. In step 710, RAMs are constructed for resources that do not already have them. The step of deploying one or more resource abstract models is illustrated in step 715, where the RAMs are distributed to their associated resources. In step 720, the configuration database is built so that application topologies can be specified. The step of using the one or more constructed models is shown in step 725. In step 725, the service level management system operates by observing service level estimates and receiving notifications from control automation. (Specification, page 14, lines 16-21).

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

(I) Whether claims 1-4, 7, 9-13, 16, 18-22, 26-28, 31 and 33 are unpatentable under 35 U.S.C. §103(a) over U.S. Patent No. 6,086,618 issued to Al-Hilali et al. (hereinafter “Al-Hilali”) in view of U.S. Patent No. 6,216,119 issued to Jannarone (hereinafter “Jannarone”).

(II) Whether claims 5, 6, 8, 14, 15, 17, 23-25, 29, 30 and 32 are unpatentable under 35 U.S.C. §103(a) over Al-Hilali in view of Jannarone in further view of U.S. Patent No. 6,959,335 issued to Hayball et al. (hereinafter “Hayball”).

ARGUMENT

Appellants incorporate by reference herein the disclosure of their previous response filed in the present application, namely the response dated January 18, 2007.

(I) Whether claims 1-4, 7, 9-13, 16, 18-22, 26-28, 31 and 33 are unpatentable under 35 U.S.C. §103(a) over U.S. Patent No. 6,086,618 issued to Al-Hilali et al. (hereinafter “Al-Hilali”) in view of U.S. Patent No. 6,216,119 issued to Jannarone (hereinafter “Jannarone”).

Regarding the §103(a) rejections, Appellants assert that the various references, alone or in combination, fail to teach or suggest all of the limitations of claim 1-33, as will be explained below. Furthermore, with regard to the combinations of the various references, Appellants assert that such combinations are improper, as will be explained below.

The Examiner cites Al-Hilali in combination with Jannarone in rejecting independent claims 1, 11, 20 and 27. More particularly, the Examiner cites portions of Al-Hilali as disclosing certain limitations of the independent claims, and cites portions of Jannarone as disclosing certain other limitations of the independent claims. Below, Appellants explain how such portions of Al-Hilali and Jannarone fail to teach or suggest what the Examiner contends that they teach or suggest. While Appellants may refer from time to time to each reference alone in describing its deficiencies, it is to be understood that such arguments are intended to point out the overall deficiency of the cited combination.

Independent claim 1 is directed to a method of constructing a model representative of a resource for use in managing a service associated with the resource, comprising the steps of associating a resource abstract model with the resource, wherein the resource abstract model is configured to automatically determine a set of resource metrics to be used to construct a model representative of the resource such that a reduced set of resource metrics is considered; and

constructing the model representative of the resource based on the reduced set of resource metrics obtained in accordance with the resource abstract model.

The Examiner in formulating the §103(a) rejection of claim 1 argues that each and every one of the above-noted limitations is met by the collective teachings of Al-Hilali and Jannarone. Appellants respectfully disagree.

In characterizing the Al-Hilali reference as allegedly meeting certain limitations of claim 1, the Examiner relies primarily on FIG. 4, column 9, lines 31-36 and column 10, lines 11-14. The Al-Hilali reference fails to disclose the limitation “wherein the resource abstract model is configured to automatically determine a set of resource metrics to be used to construct a model representative of the resource such that a reduced set of resource metrics is considered.” Rather, Al-Hilali creates a model “by first determining the resources used by the server application and defining ‘transactions’ that occur at the server application. A transaction is an identifiable operation occurring at the server application in response to user/client behavior and can typically be measured in rate form (transactions per second). Furthermore, an anticipated user load based on user behavior can also be reduced to transaction rates with a server application.” (See page 4, lines 56-64 of Al-Hilali). The Al-Hilali reference appears to be a data driven technique similar to the IBM DB2 database management system referenced above.

The Examiner looks to the Jannarone reference to supplement the above-noted deficiencies of Al-Hilali as applied to claim 1. However, the Jannarone reference also fails to teach or suggest configuring a resource abstract model “to automatically determine a set of resource metrics to be used to construct a model representative of the resource such that a reduced set of resource metrics is considered.”

In characterizing the Jannarone reference as allegedly meeting certain limitations of claim 1, the Examiner relies primarily on column 3, lines 21-34 and column 4, lines 22-32 of Jannarone. Also, in the Advisory Action, the Examiner again states that Jannarone at column 3, lines 21-34 and column 4, lines 22-32 discloses the argued limitations.

The Jannarone reference, at col. 3, lines 21-34, states the following:

The present invention meets the needs described above in a multi-kernel neural network computing architecture configured to learn correlations among feature values as the network monitors and imputes measured input values and also forecasts future output values. This computing architecture, referred to as a concurrent-learning information processor (CIP), includes a multi-kernel neural network array with the capability to learn and predict simultaneously. The CIP also includes a manager and an input-output transducer that may be used for input-output refinement. These components allow the computing capacity of the multi-kernel array to be reassigned in response to measured performance or other factors.

Jannarone, at col. 4, lines 22-32 states the following:

Generally described, the invention is a method for responding to computed output values that are based on measured input values received during a current time trial and during one or more historical time trials. The measured input values are received for the current time trial, and a vector of input feature values is assembled based on the measured input values. The input feature values are provided to a multi-kernel processor. Each kernel of the processor is operative for receiving one or more of the input feature values and performing a number of operations using the input feature values.

The relied upon portions of Jannarone do not teach or suggest the recited limitations of claim

1. No where does Jannarone disclose configuring a resource abstract model to automatically determine a set of resource metrics to be used to construct a model representative of the resource such that a reduced set of resource metrics is considered, as recited in the claimed invention. Instead, the relied-upon portions of Jannarone refers to a concurrent-learning information processor (CIP) with a multi-kernel neural network array which is capable of learning and predicting simultaneously, by responding to computed output values that are based on measured input values received from current or historical time trials, which clearly has nothing to do with configuring a resource abstract model to automatically determine a set of resource metrics to be used to construct a model representative of the resource such that a reduced set of resource metrics is considered.

Thus, the Jannarone reference fails to supplement the above-noted deficiencies of Al-Hilali as applied to claim 1. Accordingly, it is believed that the combined teachings of Al-Hilali and Jannarone fail to meet the limitations of amended claim 1.

Also, the Examiner has failed to identify a cogent motivation for combining Al-Hilali and Jannarone in the manner proposed. The Examiner provides the following statement of motivation beginning at page 4, first paragraph of the Office Action:

Al-Hilali and Jannarone are analogous art because they are both related to estimating resource usage requirements.

Therefore, it would have been obvious to one of ordinary skill in this art at the time of invention by applicant to utilize the multi-kernel neural network computing architecture of Jannarone in the method of Al-Hilali because Jannarone utilizes a multi-kernel array which learns and predicts simultaneously in “real time” (see col. 3, lines 43-45), and as a result, Jannarone reports an improvement over the drawbacks of conventional neural network systems and a significant advancement in neural network techniques (see col. 3, lines 45-48 and col. 7, lines 51-52).

The Federal Circuit has stated that when patentability turns on the question of obviousness, the obviousness determination “must be based on objective evidence of record” and that “this precedent has been reinforced in myriad decisions, and cannot be dispensed with.” In re Sang-Su Lee, 277 F.3d 1338, 1343 (Fed. Cir. 2002). Moreover, the Federal Circuit has stated that “conclusory statements” by an examiner fail to adequately address the factual question of motivation, which is material to patentability and cannot be resolved “on subjective belief and unknown authority.” Id. at 1343-1344. There has been no showing in the present §103(a) rejection of claim 1 of objective evidence of record that would motivate one skilled in the art to combine Al-Hilali and Jannarone to produce the particular limitations in question. Although the Examiner in the Advisory Action argues that “[he] elaborated a motivation for combining the references and not a motivation for combining the references ‘to come up with the invention,’” Appellants maintain that the above-quoted statement of motivation provided by the Examiner appears to be a conclusory statement of the type ruled insufficient in the In re Sang-Su Lee case. Furthermore, the proposed combination appears to be based primarily on hindsight, with the Examiner attempting to reconstruct the claimed arrangement from disparate references.

For at least these reasons, Appellants assert that claim 1 is patentable over Al-Hilali and Jannarone. Independent claims 11, 20 and 27 include limitations similar to those of claim 1, and are therefore believed patentable for reasons similar to those described above with reference to claim 1.

Furthermore, Appellants assert that the claims which depend from claim 1, 11, 20 and 27 are patentable over the Al-Hilali/Jannarone combination not only for the reasons given above with respect to claims 1, 11, 20 and 27, but also because such dependent claims recite patentable subject matter in their own right, as will be set out below.

Regarding claim 3, Al-Hilali does not disclose constructing the resource abstract model by at least one individual with expertise associated with the resource. As illustrated in FIG. 1 of the present application, a resource abstract model (RAM) associated with a resource type or instance is constructed by an expert for that resource, such as expert 105 for a type 1 resource and expert 130 for a type N resource. (Specification, page 10, lines 8-11). Instead, the relied-upon portion of Al-Hilali states that “in order to identify and define the transactions appropriate at step 102, it is helpful, though not entirely necessary, to have some ideas as to user behavior being simulated,” which is not relevant to the rejection of the claimed features in claim 3.

With regard to claims 4, 13, 22 and 28, Al-Hilali does not teach or suggest the step of obtaining one or more service level metrics for use in constructing the model representative of the resource. As noted in the present specification, management system 160 has service level agreements 165 that specify end user expectations for interactions with the enterprise system. These expectations are quantified by service level (SL) metrics 168 such as, for example, response time and throughput. Management system 160 also maintains a configuration database 170 in which the system identifies the resources that are used to deliver the services described in the service level agreement. An identification and execution engine 175 within management system 160 interacts with the resources of the enterprise system 102 to: (a) obtain the set of available metrics; and (b) collect data for these metrics. The collected data is used to construct a system model 180 that is used to report service levels. (Specification, page 10, lines 12-21).

In characterizing Al-Hilali as disclosing the limitations of claims 4, 13, 22 and 28, the Examiner refers to column 10, lines 16-20 of Al-Hilali, which states that “[b]ased on the original user interaction with the client application, the most relevant set of possible transactions may then be selected and eventually measured for resource usage leading to a more accurate and relevant model.” Al-Hilali, by selecting the most relevant set of possible transactions to be measured on the basis of

the original user interaction with the client application teaches against obtaining the one or more service level metrics, which are quantifications of end user expectations for interactions with the enterprise system.

Regarding claims 7, 16 and 31, Jannarone does not teach or suggest the resource being an element of an autonomic computing environment. Jannarone at column 16, lines 28-30 only states that “the kernel array 14 may be analogized to the autonomic processes of the thinking organism,” which is not relevant to the rejection of the claimed features in claims 7, 16 and 31.

(II) Whether claims 5, 6, 8, 14, 15, 17, 23-25, 29, 30 and 32 are unpatentable under 35 U.S.C. §103(a) over Al-Hilali in view of Jannarone in further view of U.S. Patent No. 6,959,335 issued to Hayball et al. (hereinafter “Hayball”).

Appellants assert that claims 5, 6, 8, 14, 15, 17, 23-25, 29, 30 and 32, which depend from claims 1, 11, 20 and 27, are patentable over the Al-Hilali/Jannarone/Hayball combination not only for the reasons given above with respect to claims 1, 11, 20 and 27, but also because such dependent claims recite patentable subject matter in their own right.

Regarding claims 5, 14, 23 and 29, Hayball does not teach or suggest the one or more service level metrics being obtainable from one or more service level agreements. As noted above with regard to claims 4, 13, 22 and 28, management system 160 has service level agreements 165 that specify end user expectations for interactions with the enterprise system. These expectations are quantified by service level (SL) metrics 168 such as, for example, response time and throughput. In Hayball, the service level agreement refers to guaranteeing a bandwidth and quality of service level of a path between two specified nodes. Hayball, at column 4, lines 34-38, refers to a network operator being able to meet service level agreements that are entered into and ensure that specified quality of service levels are met. The relied-upon portion of Hayball does not teach or suggest a service level metric being obtained from the one or more service level agreements.

With regard to claims 6, 15, 24 and 30, Hayball does not teach or suggest obtaining a topology of one or more resources used to deliver one or more services associated with the one or more service level agreements, including the resource for which the model is being constructed, for

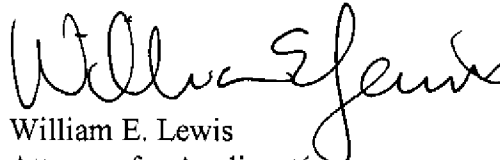
use in constructing the model representative of the resource. In characterizing Hayball as teaching the limitations of claims 6, 15, 24 and 30, the Examiner refers to column 4, lines 17-20 of Hayball, which states that “the amount of available bandwidth over said path is assessed using said model; and provisioning information to provision said path is produced using said model.” The Examiner interprets “amount of available bandwidth over said path” as the recited “topology” (Final Office Action, page 9, first paragraph). For the sake of argument, even if the amount of available bandwidth over said path is interpreted as the recited “topology,” Hayball teaches against the claimed invention by using said model to assess the amount of available bandwidth over said path, whereas claims 6, 15, 24 and 30 recite obtaining a topology of one or more resources...for use in constructing the model representative of the resource.

Regarding claims 8, 17, 25 and 32, Hayball does not disclose the constructed model useable for at least one of: (i) reporting one or more service level metrics; (ii) automating service level compliance; (iii) permitting a service provider to manage one or more service on demand; and (iv) generating one or more notifications related to automated service level enforcement. It is not clear how the relied-upon portions of Hayball, which refer to monitoring the amount of traffic issuing from particular customers, and to a condition class dependent on a certain number of times that the service level agreement is violated, teach or suggest the claimed features of claims 8, 17, 25 and 32.

It is also asserted that the motivation set forth by the Examiner to combine Hayball with Al-Hilali and Jannarone is insufficient under the In re Sang-Su Lee decision (cited above).

In view of the above, Appellants believe that claims 1-33 are in condition for allowance, and respectfully request withdrawal of the §103(a) rejections.

Respectfully submitted,

A handwritten signature in black ink, appearing to read "William E. Lewis". The signature is fluid and cursive, with the first name "William" being more prominent than the last name "Lewis".

William E. Lewis
Attorney for Applicant(s)
Reg. No. 39,274
Ryan, Mason & Lewis, LLP
90 Forest Avenue
Locust Valley, NY 11560
(516) 759-2946

Date: April 19, 2007

APPENDIX

1. A method of constructing a model representative of a resource for use in managing a service associated with the resource, comprising the steps of:

associating a resource abstract model with the resource, wherein the resource abstract model is configured to automatically determine a set of resource metrics to be used to construct a model representative of the resource such that a reduced set of resource metrics is considered; and

constructing the model representative of the resource based on the reduced set of resource metrics obtained in accordance with the resource abstract model.

2. The method of claim 1, wherein the constructed model comprises a quantitative model.

3. The method of claim 1, wherein the resource abstract model is constructed by at least one individual with expertise associated with the resource.

4. The method of claim 1, further comprising the step of obtaining one or more service level metrics for use in constructing the model representative of the resource.

5. The method of claim 4, wherein the one or more service level metrics are obtainable from one or more service level agreements.

6. The method of claim 5, further comprising the step of obtaining a topology of one or more resources used to deliver one or more services associated with the one or more service level agreements, including the resource for which the model is being constructed, for use in constructing the model representative of the resource.

7. The method of claim 1, wherein the resource is an element of an autonomic computing environment.

8. The method of claim 1, wherein the constructed model is useable for at least one of: (i) reporting one or more service level metrics; (ii) automating service level compliance; (iii) permitting a service provider to manage one or more service on demand; and (iv) generating one or more notifications related to automated service level enforcement.

9. The method of claim 1, further comprising the step of checking the accuracy of the constructed model.

10. The method of claim 9, wherein the accuracy checking step comprises use of change point detection.

11. Apparatus for constructing a model representative of a resource for use in managing a service associated with the resource, comprising:

a memory; and

at least one processor coupled to the memory and operative to: (i) automatically determining, via a resource abstract model, a set of resource metrics to be used to construct a model representative of the resource such that a reduced set of resource metrics is considered; and (ii) constructing the model representative of the resource based on the reduced set of resource metrics obtained in accordance with the resource abstract model.

12. The apparatus of claim 11, wherein the constructed model comprises a quantitative model.

13. The apparatus of claim 11, wherein the at least one processor is further operative to obtain one or more service level metrics for use in constructing the model representative of the resource.

14. The apparatus of claim 13, wherein the one or more service level metrics are obtainable from one or more service level agreements.

15. The apparatus of claim 14, wherein the at least one processor is further operative to obtain a topology of one or more resources used to deliver one or more services associated with the one or more service level agreements, including the resource for which the model is being constructed, for use in constructing the model representative of the resource.

16. The apparatus of claim 11, wherein the resource is an element of an autonomic computing environment.

17. The apparatus of claim 11, wherein the constructed model is useable for at least one of: (i) reporting one or more service level metrics; (ii) automating service level compliance; (iii) permitting a service provider to manage one or more service on demand; and (iv) generating one or more notifications related to automated service level enforcement.

18. The apparatus of claim 11, wherein the at least one processor is further operative to check the accuracy of the constructed model.

19. The apparatus of claim 18, wherein the accuracy checking operation comprises use of change point detection.

20. An article of manufacture for constructing a model representative of a resource for use in managing a service associated with the resource, comprising a machine readable medium containing one or more programs which when executed implement the steps of:

automatically determining, via a resource abstract model, a set of resource metrics to be used to construct a model representative of the resource such that a reduced set of resource metrics is considered; and

constructing the model representative of the resource based on the reduced set of resource metrics obtained in accordance with the resource abstract model.

21. The article of claim 20, wherein the constructed model comprises a quantitative model.

22. The article of claim 20, further comprising the step of obtaining one or more service level metrics for use in constructing the model representative of the resource.

23. The article of claim 22, wherein the one or more service level metrics are obtainable from one or more service level agreements.

24. The article of claim 20, further comprising the step of obtaining a topology of one or more resources used to deliver one or more services associated with the one or more service level agreements, including the resource for which the model is being constructed, for use in constructing the model representative of the resource.

25. The article of claim 20, wherein the constructed model is useable for at least one of: (i) reporting one or more service level metrics; (ii) automating service level compliance; (iii) permitting a service provider to manage one or more service on demand; and (iv) generating one or more notifications related to automated service level enforcement.

26. The article of claim 20, further comprising the step of checking the accuracy of the constructed model.

27. A method of providing resource management services, comprising the steps of:
deploying one or more resource abstract models in association with one or more resources,
wherein each of the one or more resource abstract models is configured to automatically determine a

set of resource metrics to be used to construct a model representative of the resource such that a reduced set of resource metrics is considered; and

based on the one or more reduced sets of resource metrics obtained in accordance with the one or more resource abstract models, constructing one or more models representative of the one or more resources; and

using the one or more constructed models to manage the one or more resources.

28. The method of claim 27, further comprising the step of obtaining one or more service level metrics for use in constructing the one or more models representative of the one or more resources.

29. The method of claim 28, wherein the one or more service level metrics are obtainable from one or more service level agreements.

30. The method of claim 27, further comprising the step of obtaining a topology of the one or more resources used to deliver one or more services associated with the one or more service level agreements, for use in constructing the one or more models representative of the resource.

31. The method of claim 27, wherein the resource is an element of an autonomic computing environment.

32. The method of claim 27, wherein the one or more constructed models are useable for at least one of: (i) reporting one or more service level metrics; (ii) automating service level compliance; (iii) permitting a service provider to manage one or more service on demand; and (iv) generating one or more notifications related to automated service level enforcement.

33. The method of claim 27, further comprising the step of checking the accuracy of the one or more constructed models.

EVIDENCE APPENDIX

None.

RELATED PROCEEDINGS APPENDIX

None.